arXiv:1612.04901v1 [cs.CV] 15 Dec 2016

Tinkering Under the Hood: Interactive Zero-Shot Learning with Net Surgery

Vivek Krishnan Deva Ramanan
Carnegie Mellon University
Pittsburgh, PA 15213
{vrkrishn,deva}@cs.cmu.edu

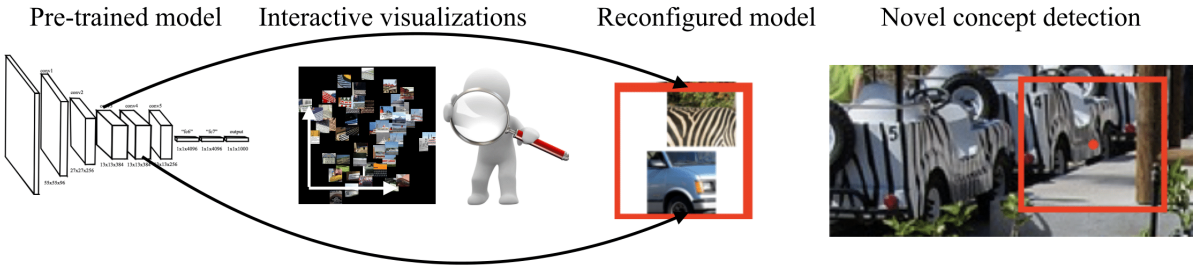

We introduce a pictorial language for visualizing, understanding, and *re-wiring* internal representations learned by CNNs to detect never-before-seen concepts, such as a "striped car".

Abstract

We consider the task of visual net surgery, in which a CNN can be reconfigured without extra data to recognize novel concepts that may be omitted from the training set. While most prior work make use of linguistic cues for such "zero-shot" learning, we do so by using a pictorial language representation of the training set, implicitly learned by a CNN, to generalize to new classes. To this end, we introduce a set of visualization techniques that better reveal the activation patterns and relations between groups of CNN filters. We next demonstrate that knowledge of pictorial languages can be used to rewire certain CNN neurons into a part model, which we call a pictorial language classifier. We demonstrate the robustness of simple PLCs by applying them in a weakly supervised manner: labeling unlabeled concepts for visual classes present in the training data. Specifically we show that a PLC built on top of a CNN trained for ImageNet classification can localize humans in Graz-02 and determine the pose of birds in PASCAL-VOC without extra labeled data or additional training. We then apply PLCs in an interactive zero-shot manner, demonstrating that pictorial languages are expressive enough to detect a set of novel visual classes in MS-COCO.

1. Introduction

Convolutional Neural Networks (CNNs) have revolutionized modern vision pipelines by replacing hand-crafted features with large set of trainable parameters that learn complex image representations from data. However, these models notoriously require a substantial amount of annotated training data to learn a feature hierarchy that generalizes well. As a result, subsequent performance gains in various visual tasks have been limited by the availability of large datasets such as ImageNet [4]

Explainable AI: A growing reactionairy response to deep networks is that their opaqueness makes them difficult to use when more than a simple prediction is required. If a network fails, why? How will a network behave on never-before-seen data? To answer such questions, there is a renewed interest in so-called *explainable AI*[1]. One goal in this reinvigorated field is the development of machine learning systems that are designed to be more interpretable and explanatory.

Interpreting Networks: Rather than redesigning a deep network architecture to become more interpretable, a complementary approach is to *post-hoc* add interpretable se-

[1] http://www.darpa.mil/program/explainable-artificial-intelligence

mantics to an existing model. Deep hierarchical models often dramatically outperform manually-designed hierarchies (classically based upon edges, parts, objects [18]), so it would be valuable to understand what non-obvious concepts are being learned by such high-performing deep networks. An illustrative example might be post-hoc interpretation of the AlphaGo network [25]; while an automated Go player itself is certainly interesting, what maybe more so is understanding the non-obvious heuristics it learns to employ.

Model reconfigurability: Evaluating post-hoc interpretability is challenging. In this work, we propose an approach based on reconfigurability; that is, we quantitatively evaluate the interpretability of a model by the *ability to interactively re-purpose it for never-before-seen tasks*. A practical issue with virtually any real-world deployment of machine-learned models is the ability to adapt to novel conditions. Such a problem formulation is often cast as domain adaption or transfer learning. Indeed, virtually all state-of-the-art vision systems make use of transfer learning via models pre-trained on ImageNet [9, 28]. A related but distinct formalism is that of domain adaptation, where a pre-trained model is adapted to a novel domain with different statistics from test data [1, 24] or entirely different modalities [11, 26]. However, in most systems, some amount of target data for the task of interest is needed to provide a gradient-based signal for learning.

Interactive zero-shot learning: The most difficult instantiation may arguably be the zero-shot setting, where no data for the novel task is provided. In this setting, even specifying the new task is difficult. One traditionally makes use of an external semantic knowledge base to do so [22]. For example, an existing lexical database may define an "elund" as a "striped animal with horns". This knowledge can be combined with existing striped and horn models to yield a visual model of a never-before-seen object. The vast majority of past work makes use of linguistic semantics, mined from text corpora [7] or lexical knowledge bases such as Wordnet [20]. Instead, we propose to use a human as the external knowledge base, yielding a novel form of *interactive zero-shot learning*. The output of such a process is a model where spatial logic and combination of parts follows directly from human logic.

Pictography: Rather than relying on linguistics, our approach encodes semantic concepts pictorially in a manner that is intuitive for a human user. Cognitive psychology suggests that many concepts in one's associative memory may be better represented through visual imagery rather than linguistic tokens [21]. Similarly, it is well-known in developmental psychology that children learn spatial concepts visually rather than linguistically [3]. Inspired by these observations, we demonstrate that CNN visualizations can be used to define a *pictorial language* that can be used to interactively reconfigure networks for novel tasks.

Visualizations: Our pictorial approach builds on a body of work that aims to visualize and understand internal representations in CNNs. Common visualization techniques include reconstructing image filters (by deconvolution [29]), inverting feature transformations [17], or computing the set of image patches that maximally activate particular neurons [9, 30]. Instead, we interpret local neighborhoods of neural activations as *embeddings*, and use classic techniques for visualizing low-dimensional projections such as tSNE [27, 16]. Importantly, our visualizations allow users to discover new pictorial concepts that cannot be represented by any single neuron. We then introduce a simple visual grammar for composing these pictorial concepts together to describe never-before-seen objects (such as a "striped car"). Our grammars allow a user to build zero-shot classifiers composed out of visual tokens that may not have a natural linguistic "name". Our grammars can be seen as spatial part models [10, 31] where part responses are equivalent to interactively-defined functions of activations (computed through our embedding visualizations).

Overview: We begin with a novel techniques for visualizing neurons in deep networks based on embeddings (Sec. 2). We describe a pictorial grammar that uses such visualizations to repurpose existing deep networks for novel tasks (Sec. 3). We use these techniques in Sec. 4 to repurpose a network trained for image *classification* ([13]) for object *detection*, including never-before-seen objects.

2. Visualizing Activations

In this section, we describe our approach for visualizing and understanding CNNs. Though our methods directly apply to any existing network, we illustrate them on our running example of AlexNet. Our methods fundamentally require both a pre-trained CNN and a collection of validation images with which to process the CNN and diagnose its behaviour. We use the ILSVRC 2012 validation set of 50,000 images [23]. Importantly, this set is distinct from the test images we use for evaluation.

2.1. Retinotopic embeddings

The heart of our approach requires the ability to interpret the semantics of particular neurons, where neurons correspond to individual activations computed throughout the layers of a CNN. Note that it is not clear that this is even possible; indeed, the well-known "grandmother neuron" hypothesis that individual neurons represent semantic concepts (such as one's grandmother) seems at odds with notion that deep networks are distributed (rather than localized) representations [2]. We will shortly offer a solution to this apparent contradiction.

Following established convention, we will denote activations by their layer, such as `conv3`, `conv4`, etc. While

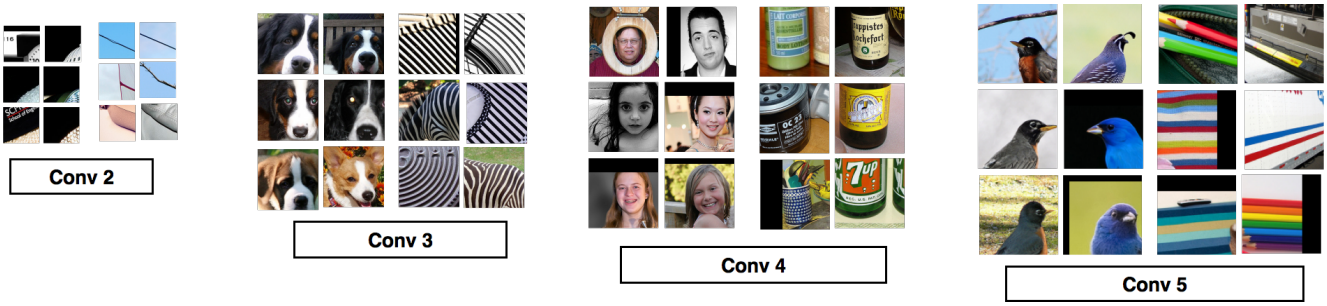

Figure 1. Visualization of image features from AlexNet. For layers 2-6, we show the top 6 activations taken from the ILSVRC 2012 [23] validation set. The size of the image patch corresponds to the theoretical RF of the neuron as computed in [15, 8]. We see that (1) a CNN trained for object classification learns a rich, hierarchical set of features, ranging from textures to complex invariances such as pose. (2) Neurons can learn to identify and localize unlabeled concepts (such as face). (3) These images of faces belong to a wide range of classes including Helmet, French Horn, Suit, Bathtub, and Park Bench. Even though there is no category on ILSVRC 2012 for generic human detection, the CNN has automatically determined that it is useful to reason about human faces as intermediate concepts.

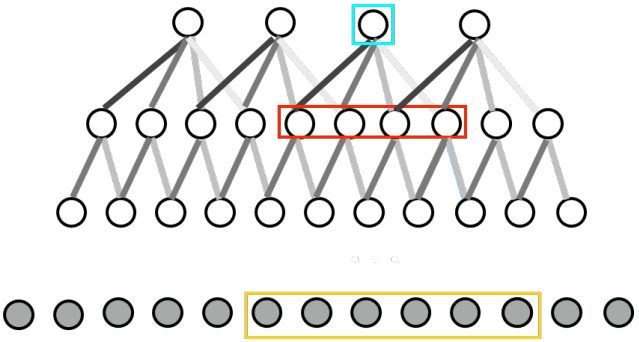

Figure 2. Our retinotopic view of CNNs. Here a is the activation of particular CNN neuron (cyan). We embed this neuron as a feature vector $w \cdot x + b$, where x is a set of input activations (red). The receptive field in the original image corresponding to neuron a is drawn in yellow.

numerous techniques have been proposed for understanding the semantics of such neurons, our starting point is a surprisingly simple strategy: characterize a neuron by the *set* of N image patches that maximally activate that particular neuron. Such an approach was popularized by [9, 30]. We show examples for $N = 6$ for a variety of neurons across different layers in Fig. 1; one can readily identify neurons corresponding to different textures, parts, and objects. Neurons in lower layers are characterized by smaller *receptive fields* (defined as the set of pixel values that will affect a given neural activation). Neurons in deeper layers tend to be invariant to various factors such as pose, but perhaps surprisingly, strong semantics appear even at shallower layers (e.g., the dog-face neuron in `conv3`).

Embeddings: Because neurons at higher layers tend to be more invariant, we would like to understand the *variation* in the appearance of image patches that activate a particular neuron. Recall that a neuron activation a is computed from previous layer responses with a (convolutional) linear dot product with weights w and bias b, followed by a (ReLU) nonlinearity:

$$a = \max(0, w \cdot x + b), \tag{1}$$

where x corresponds to a local neighborhood of activations from the previous layer. For concreteness, if a was from `conv3`, x is of dimensionality $3 \times 3 \times 256$. We propose to view $x \in \mathcal{X}$ as a point in a $\mathcal{X} = R^{2304}$ dimensional embedding space, where the activation a is given by a linear threshold function in $\mathcal{X}$. We posit that this embedding space is a rich characterization of the invariances learned by this particular neuron - e.g., if neuron a corresponds to a view-invariant object, then different points in $\mathcal{X}$ may correspond to different viewpoints of that object. We use the embedding space $\mathcal{X}$ to visualize image patches of size given by the receptive feild of neuron a (Fig. 2 and Fig. 3).

Past work: The idea of viewing neural activations as embeddings is certainly not new. Training deep nets to predict embedding coordinates is a well-studied task [12]. Moreoever, it is widely known that one can interpret the final (FC7) activation layer as an embedding, even for nets originally trained for classification [5]. Our key insight simply extends this observation to convolutional layers, which we demonstrate requires examining *local neighborhoods* of activations (1). This observation might offer a solution to the localized (grandmother) versus distributed debate: interpretable semantics might reside in localized neighborhoods of activations rather than any individual one. Or in other terms, we should **not** conclude that a network cannot understand "wheels" simply because we cannot find any single neuron that responds to "wheels": rather, this knowledge may be encoded in a distributed fashion across a local neighborhood of neurons. Because such neighborhoods share similar receptive fields (due to the convolutional structure of weights), we refer to such embeddings as *retinotopic embeddings*. This characterization of deep nets as learning progressively invariant embeddings may agree with per-

spectives of such models as "disentangling" factors of variation from one layer to the next [2].

Dimensionality reduction: We use the above observation to visualize N high-scoring patches for a particular neuron. By viewing each patch as a point in $\mathcal{X}$, we can visualize a collection using low-dimensional projections (Fig. 3). T-SNE [27] is a popular technique for visualizing *nonlinear* projections that preserves neighborhood structure in the high-dimensional space. We also experimented with simpler *linear* projections obtained through PCA, which also allows us to visualize neuron a as a 2D line given by the linear threshold function (w, b). Formally speaking, let us write the 2-dimensional projected coordinates of a point $x \in \mathcal{X}$ as

$$(c_1, c_2) = (v_1 \cdot x, v_2 \cdot x) \quad \text{[PCA]},$$

where (v_1, v_2) are projection vectors given by PCA. Given that we can approximately reconstruct x from its projection with $x \approx c_1 v_1 + c_2 v_2$, the linear threshold function from (1) can be written as

$$wx + b \approx c_1 v_1 \cdot w + c_2 v_2 \cdot w + b,$$

or a 2D line with a normal vector $(v_1 \cdot w, v_2 \cdot w)$ and offset b. This line is visualized in blue in Fig. 3.

Activations as 1-D embeddings: It is worth noting that the *un*thresholded activation value $a' = w \cdot x + b$ is itself a one-dimensional linear projection of the embedding space $\mathcal{X}$. Hence prior work that analyzes networks by plotting the top N scoring patches that activate any neuron [9, 30] can also be thought to employ low-dimensional (or rather, one-dimensional) embeddings. By projecting to more than a single dimension, our approach captures much richer semantics such as pose and categorical variation (Fig. 3).

User-drawn concepts (PCA): Our embedding visualizations can be used to define *new* linear threshold boundaries, corresponding to user-defined visual concepts. In the case of PCA-embeddings, these user-drawn linear thresholds can be back projected to define new filters and biases. Formally, given a user-drawn line with normal vector (α_1, α_2) and offset β, the corresponding high-dimensional filter producing the same response is given by

$$\begin{aligned} \alpha_1 c_1 + \alpha_2 c_2 + \beta &= \alpha_1 v_1 \cdot x + \alpha_2 v_2 \cdot x + \beta \\ &= w' \cdot x + b', \end{aligned}$$

where $w' = \alpha_1 v_1 + \alpha_2 v_2$ and $b' = \beta$. We will show examples of zero-shot models built with user-defined filters in our experimental results.

User-drawn concepts (TSNE): Reconstructing the filter from a TSNE embedding is much more difficult. TSNE does not explicitly compute an embedding function, but rather directly outputs an embedding of a fixed set of input points (implying that the embedding cannot be applied to "out-of-sample" points). P-TSNE [16] is a parametric extension that essentially trains a feedforward neural net to predict the TSNE embedding obtained for a fixed set of inputs - e.g.,

$$(c_1, c_2) = (f_1(x), f_2(x)) \quad \text{[TNSE]}.$$

Once this function is learned, it can be applied to new "out-of-sample" patches. Because the function is nonlinear, halfspaces in the embedded space do not correspond to simple linear filters in $\mathcal{X}$. However, one can still "implement" a user-drawn halfspace by explicitly computing the feedforward embedding on any query patch, and linearly scoring the resulting 2D projection:

$$a_{TSNE} = \max(\alpha_1 f_1(x) + \alpha_2 f_2(x) + \beta, 0).$$

Note that this same approach could also apply for PCA; instead of reconstructing the original filter, one could explicitly computes linear thresholds in the low-dimensional projected space

$$a_{PCA} = \max(\alpha_1 v_1 \cdot x + \alpha_2 v_2 \cdot x + \beta, 0),$$

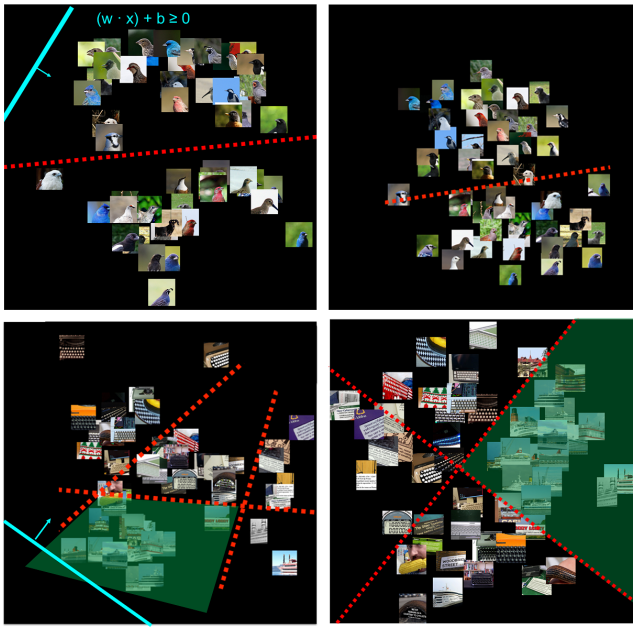

Figure 3. We visualize pairs of (PCA,TSNE) embedding of two `conv5` neurons, along with user-drawn linear boundaries (dotted red lines. For the PCA visualizations (left), we show a projection of the original filter weight into 2D space using the PCA basis (blue line). Because the t-SNE projections (right) are non-linear, we cannot simply project the filter weight in such a fashion. For the **top** pair, we clearly observe a linear separator in the embeddings between bird heads with different orientations. The **bottom** illustrates a grouping of texture patterns that may not be easily described linguistically, but appear to loosely correspond to keyboard, cruise ship, and text. We have highlighted the region corresponding to cruise ship as an example.

Figure 4. We illustrate two search methods based on neuron "58" from `conv5`, which fires on images of vehicles. We provide the 4 image patches in the ILSVRC validation set that maximally activate the neuron (left). Weight Search (**left**) consists of finding the maximum weight for each spatial location in the filter. We plot the learned weight values for each filter and indicate the maximum filter weight with a red line. Input Co-occurence Search (**right**) consists of averaging the input features over the top-100 activations of the neuron. Here we plot the averaged input for each filter and highlight the maximum with a red line. For each search method, we also show the top image at each spatial location corresponding to the retrieved `conv4` filters.

which might be faster if processing multiple user-defined concepts obtained from the same embedding.

2.2. Searching for semantics

The previous section describes an approach for uncovering the semantic concepts of individual neurons, as well as constructing new user-defined concepts. In theory, one could apply the given approach in an exhaustive fashion over all neighborhoods of neurons to uncover semantics of interest, but methods to filter the set of neurons can be used to reduce the search time.

We envision two modes of user interaction for our zero-shot pipeline. A user may have a semantic concept in mind ("I want to build a model of 3-wheeled vehicles"), or users may wish to peruse internal concepts captured in a network in an exploratory fashion. In either case, the user will identify a number of visual concepts which will then be assembled into a pictorial grammar (Sec. 3). In order to identify a pool of relevant candidates, a successful pipeline appeared to be first identifying an "interesting" neuron (by random search at layers in `conv5` or higher), and then searching for related neurons as detailed below. For example, after finding a "face" neuron, one might then search for neurons that co-activate, which might correspond to other body parts. Key to these methods is that statistics about the network and training set can be precomputed, making subsequent queries on the network available in constant time.

Filter search: The first method utilizes the filter weights learned by the CNN. Recall that activations are computed by linear threshold functions (1) where the inputs x themselves must be non-negative (since they are rectified activations themselves). This implies that the magnitude of weight of each input is directly related to the final activation value of a neuron, as well as its importance in positive feature detection for the neuron. Therefore, given a neuron of interest a, we search for connected neurons at higher and lower levels associated with large weights w.

Co-occurrence search: Our second method empirically searches for collections of neurons that consistently activate (over a set of images). Given a particular neuron of interest, we compute all non-zero activations on a validation set, and record the average activation of all connected neurons in higher and lower layers. We construct a pool of connected neurons (ordered by their filter weight or empirical correlation) that serve as candidates for embedding visualizations. We show examples of successfully-identified related neurons in Fig. 4.

3. Pictorial Language Classifiers

Our previous sections describe an interface for identifying visual tokens of interest. Examples might include textures ("striped"), part decompositions ("wheels"), and object invariances ("frontal-view faces"), that together can represent a wide range of complex objects. In this section, we construct a simple visual grammar for composing these tokens into flexible models for object detection. Importantly, the final visual grammar can be executed using standard CNNs (convolution and max-pooling) implying that the user-constructed models can be efficiently implemented as additional add-on layers to AlexNet. Since the vast majority of computation will be shared, this means that user-defined object detection essentially comes "for-free".

Our object detection model is based off the deformable parts model (DPMs) described in [6]. In particular, we exploit the insight in [10] that DPMs can be implemented in CNN toolboxes. We revisit the derivation with slightly more detail here, pointing out that general tree-structured grammars can be implemented in CNNs. This allows, for example, users to build zero-shot detectors for objects (e.g., a "striped car") as well as spatial arrangements of objects (e.g., "a striped car next to two red bicycles"). Let us write

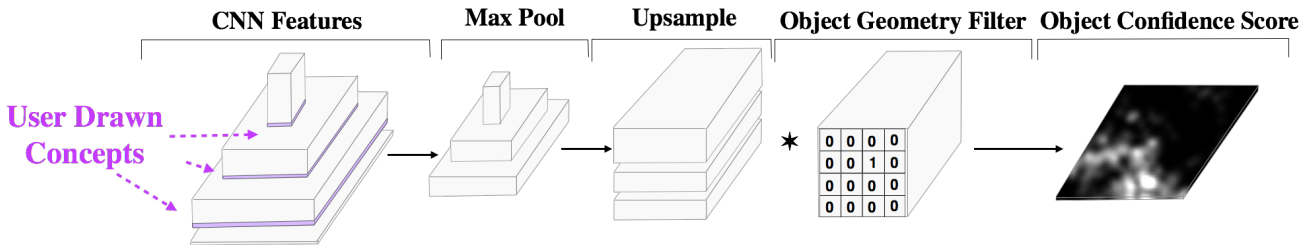

Figure 5. We efficiently implement pictorial grammars by augmenting pre-trained CNNs with additional user-defined concept filters (from Sec. 2), max-pooling responses, upsampling responses across different layers to the same pixel resolution, and summing max-pooled response maps shifted by the anchor location of a part with respect to its parent (implemented as a sparse object geometry filter [10]).

the pixel location of part i as $z_i = (x_i, y_i)$, and the score of a collection of parts $z = \{z_i\}$ as follows:

$$\text{score}(z) = \sum_i \phi(z_i) + \sum_{j \in parent(i)} \psi(z_j - z_i - a_j) \quad (2)$$

The local function $\phi(z_i)$ denotes the score associated with placing part i at pixel location i, which is found by evaluating (a possibly-upsampled) activation heatmap. The pairwise function $\psi(z_j - z_i - a_j)$ encodes a spatial model that denotes valid relative locations of part j to its parent i, where its rest anchor location is given by a_j. We use a simplified variant where

$$\psi(z_j - z_i - a_j) = \begin{cases} 0, & \text{if} \quad ||z_j - z_i - a_j||_\infty \leq r_j \\ -\infty, & \text{otherwise} \end{cases} \quad (3)$$

which requires parts to lie within a square neighborhood (of width $2r_j$) of their anchor location. It is well-known that the best-scoring configuration of parts can be found with dynamic programming for tree-structured spatial constraints. By using the particular spatial model given above, partial scores of each part can be efficiently computed by repeating the following updates from the leaf to the root:

$$\text{score}_i(z_i) = \phi(z_i) + \sum_{j \in kids(i)} m_j(z_i + a_j),$$

$$\text{where} \quad m_j(z_i) = \max_{\{z_j : ||z_j - z_i||_\infty \leq r_j\}} \text{score}_j(z_j)$$

Here $\text{score}_i(z_i)$ is initialized to $\phi(z_i)$ for leaf parts and $\text{score}_i(z_i)$ represents the true max-marginal score for the root part. Note that $m_j(z_i)$ is computed by max-pooling while $\text{score}_i(z_j)$ is computed by summing up shifted response maps (which can be implemented with a sparse, binary convolution operation). This implies that tree-structured dynamic programming operations can be implemented as standard layers in a CNN (Fig. 5).

User-defined spatial models: Though a tree-structured formulation is flexible, it may be noninuitive for a user to manually specify. For that reason, we limit our results to "star" structured models that consist of a single root filter and a collection of child filters. We use an empty root filter of a fixed size $H \times W$, which specifies the bounding box dimension of a candidate detection. We allow users to specify the a_j and r_j, which corresponds to the "rest" location of part within the bounding box, and the amount of valid displacements about that location. To find objects of different sizes, we process an image pyramid.

4. Evaluation

We evaluate our pictorial zero-shot framework with a variety of tasks of increasing difficulty, in terms of novelty with respect to AlexNet's training set. We first learn a model of categories that do appear in the training set but are *not labeled*. We then learn a *fine-grained* subcategory model for labeled categories, and finally conclude with a true zero-shot model of *never-before-seen* categories.

Unlabeled category (Graz pedestrians): Perhaps surprisingly, pedestrians are not an explicit category in the ILSVRC12's training set. We use our interface to construct a simple model for pedestrians using three user-defined parts, roughly corresponding to face, lower-body, and upper-body (visualized in Fig. 6). We benchmark our model as a pedestrian detector on the Graz-02 [19] dataset, which is composed of natural-scene images that contain cluttered backgrounds and complex objects in a variety of poses. We use the standard detection metric of average precision (AP) computed over various intersection-over-union (IOU) thresholds in Fig. 7. We refer the reader to the caption for more details, but in summary, star-structured part models significantly outperform any single part, as well as a "bag-of-parts" model without any spatial constraints. Our zero-shot models don't quite match the previously published performance of supervised models. One reason is that our models are tuned for frontal people. But our performance is impressive in that we use **zero** labeled examples and **no** linguistic knowledge!

Unlabeled subcategory (PASCAL birds): We next learn novel subcategory for the labeled category of birds. We specifically use the user-drawn boundary in Fig. 3-(a) to define left-facing and right-facing bird models. We evaluate

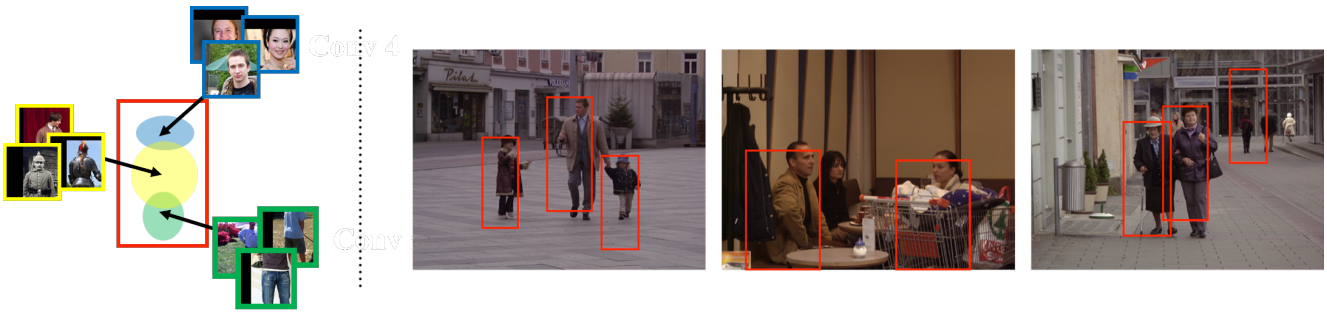

Figure 6. On the **left**, we show a visual representation of the parts and spatial constraints used to construct our PLC for human detection. For each part, we draw an arrow to the spatial anchor in the template and visualize the allowed deformation of the part as a colored region. In practice, we use max pooling to enforce this deformation and thus the region will have a rectangular shape. At **right**, we show the top scoring detections for 3 instances of Graz-02. Quantitative results are included in Fig. 7.

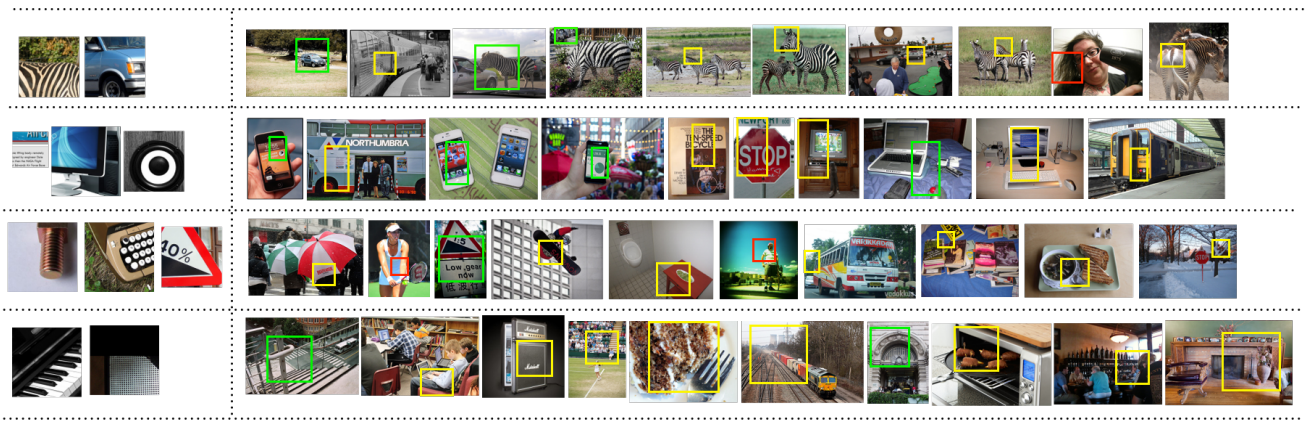

Figure 9. Concept discovery on MSCOCO. Each row presents a set of neurons that composed into a "bag-of-parts" PLC, followed by the top 10 detections from MSCOCO-val. For reference, the visual concepts were originally selected by a user wishing to model a "striped car", a "cell phone", a "geometric shape with texture", and a purely random combination of neurons. We evaluate results with a user-study where participants (none of whom designed the queries) are prompted with the following: `Interpret the query on the left as a set of visual concepts such as textures, parts, and objects. For each image on the right, determine whether the detection enclosed by the bounding box (i) represents all the queried visual concepts, (ii) represents at least one, or (iii) does not represent any.` Boxes are colored green (i), yellow (ii), or red (iii) based on the majority response. Precision for the four models at a recall of 10 is 0.401 for (i) and 0.931 for (ii), indicating that fairly accurate retrieval is possible even in the zero-shot setting.

Figure 10. We explore the notion of negative parts, by subtracting an *animal head* part score from a *stripes+car front*. On the **top**, we show three images corresponding to the query from the first row of Fig. 9, followed by score maps of the original two-part model and the two-parts-minus-the-head. Scores associated with zebra regions decrease in magnitude, but the striped car detection (on the **top-right**) remains high. This translates to a better re-ranking of retrieved images (**bottom**).

our two pose-specific bird models on Pascal VOC 2007 test-set, and show the top-3 hits for each model in Fig. 8. When we evaluate pose accuracy on the set of detected birds in PASCAL 2007-val, we obtain an accuracy of 87.5%. These results empirically support the claim that CNNs trained for image classification implicitly learn pose invariance, but

Model	AP (0.3)	AP (0.5)	AP (0.8)
PLC-Face	0.152	0.123	0.0015
PLC-Lower Body	0.412	0.036	0.006
PLC-Upper Body	0.110	0.099	0.012
PLC-All (Bag)	0.478	0.223	0.021
PLC-All (Spatial)	0.696	0.499	0.110
DPM [6]	-	0.880	-
Fast R-CNN [8]	-	0.882	-

Figure 7. Evaluation of a human detector on Graz-O2. For each image, we run the detector (which is equivalent to running Alexnet convolutionally) at 5 scales and report the highest-scoring bounding boxes after NMS. We explore variants of a PLC with 3 parts, including individual parts, a bag-of-parts (without a spatial term), and a full 3-part model with star-spatial structure (which outperforms all variants). For comparison, we show average precision (AP) results of two supervised methods at varying overlap thresholds of .3, .5, and .8. While not state-of-the-art, our results are impressive given that no person labels were ever used. Our results are particularly accurate at lower overlap thresholds. Qualitative results are included in Fig. 6.

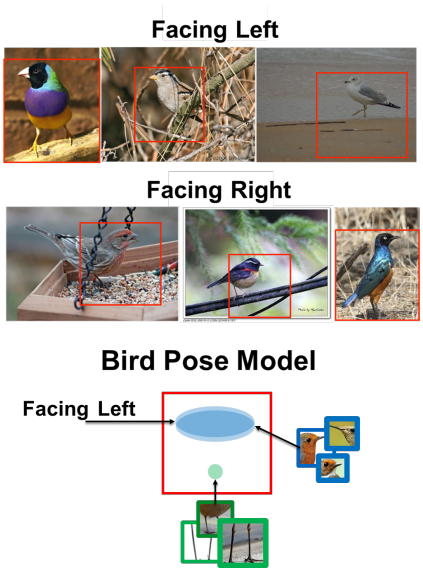

Figure 8. We apply two *pose-specific* PLCs composed of 4 parts to detect birds and estimate their pose on PASCAL VOC 2007, showing the top-3 results for each detector. As in the previous visualization, we show both the anchor locations and allowable deformation for each set of parts. The part labeled *facing left* refers to the bird pose neuron described in the text. When we evaluate pose accuracy on the set of detected birds in PASCAL 2007-val, we obtain a correct classification rate of 87.5% (where chance is roughly 50%).

this knowledge might be "tangled" in internal neural activations. Our pictorial visualizations allow a user to "untangle" such cues and explicitly build fine-grained detectors **without** any labeled examples.

Discovering concepts (MS COCO): We now use zero-shot visualizations to discover novel concepts in MSCOCO [14]. Fig. 9 shows 4 visual concept queries and the top 10 detections for each. For reference, the visual concepts were originally selected by a user wishing to model a "striped car", a "cell phone", a "geometric shape with texture", and a purely random combination of neurons. Each visual concept is user-defined to be a "bag of (2-3) parts". To evaluate results, we ask a set of participants whether or not the returned detections (i) encompass all of the visual parts, (ii) contain at least one of the visual parts, or (iii) are false detections. Our mAP for the four models at a recall of 10 is 0.401 for (i) and 0.931 for (ii). User-defined models tend to fire on image regions with a single strong part activation. We posit that since these parts rarely (if ever) co-occured in the training set, their activations are not properly calibrated, implying that better normalization of their activations might help. One interesting error mode is the *striped car* query (first row), which actually fires on a striped car (the ranked-4 result) but also finds zebras. During our semantic exploration, it is quite apparent that an *animal head*s co-occur with *stripes*. We built a PLC with a "negative" part simply by subtracting the score of the best *animal head* within a candidate bounding box. We find that this leads to noticeably better score maps and retrievals (Fig. 10), suggesting that negative parts might be a natural extension for making our interactive grammar more flexible.

Conclusion: We have shown that a CNN trained for ImageNet classification learns an image representation that can be transferred to build part-based detectors for (both previously seen but unlabeled and never-before-seen) objects without additional training data. To do so, we introduce the problem of *interactive zero-shot learning*, where a human-in-the-loop both (1) interprets and (2) reconfigures the internal semantics of a pre-trained neural net for such novel tasks. To (1) interpret internal semantics, we propose a novel perspective of convolutional neural nets as hierarchical compositions of *retinotopic embeddings*. We show that embedding visualizations can be used to interactively uncover semantic concepts that may not be captured in any single neuron. To (2) reconfigure these semantic concepts, we introduce a spatial grammar of *pictorial language classifiers* that make use of pictorial concepts rather than linguistic ones. Conveniently, pictorial language classifiers can be naturally integrated as "add-on" layers on top of existing networks.

References

[1] S. Ben-David, J. Blitzer, K. Crammer, F. Pereira, et al. Analysis of representations for domain adaptation. *NIPS*, 19:137, 2007.

[2] Y. Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.

[3] M. Bowerman. The origins of childrens spatial semantic categories: Cognitive versus linguistic determinants. *Rethinking linguistic relativity*, pages 145–176, 1996.

[4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009.

[5] J. Devlin, S. Gupta, R. Girshick, M. Mitchell, and C. L. Zitnick. Exploring nearest neighbor approaches for image captioning. *arXiv preprint arXiv:1505.04467*, 2015.

[6] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *PAMI*, 32(9):1627–1645, 2010.

[7] Z. Fu, T. Xiang, E. Kodirov, and S. Gong. Zero-shot object recognition by semantic manifold distance. In *CVPR*, pages 2635–2644, 2015.

[8] R. Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015.

[9] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014.

[10] R. Girshick, F. Iandola, T. Darrell, and J. Malik. Deformable part models are convolutional neural networks. In *CVPR*, pages 437–446, 2015.

[11] X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, pages 513–520, 2011.

[12] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

[13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.

[14] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*. 2014.

[15] J. L. Long, N. Zhang, and T. Darrell. Do convnets learn correspondence? In *Advances in Neural Information Processing Systems*, pages 1601–1609, 2014.

[16] L. Maaten. Learning a parametric embedding by preserving local structure. In *International Conference on Artificial Intelligence and Statistics*, pages 384–391, 2009.

[17] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *CVPR*, pages 5188–5196. IEEE, 2015.

[18] D. Marr. Vision: A computational investigation into the human representation and processing of visual information. 1982.

[19] M. Marszatek and C. Schmid. Accurate object localization with shape masks. In *CVPR*, pages 1–8. IEEE, 2007.

[20] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 1995.

[21] A. Paivio. Mental imagery in associative learning and memory. *Psychological review*, 1969.

[22] M. Palatucci, D. Pomerleau, G. E. Hinton, and T. M. Mitchell. Zero-shot learning with semantic output codes. In *NIPS*. 2009.

[23] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[24] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226. Springer, 2010.

[25] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[26] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks. In *ICCV*, pages 4068–4076, 2015.

[27] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.

[28] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In *NIPS*, pages 3320–3328, 2014.

[29] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Computer vision–ECCV*, pages 818–833. Springer, 2014.

[30] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014.

[31] S.-C. Zhu and D. Mumford. *A stochastic grammar of images*. Now Publishers Inc, 2007.